

Reimagining WLCG for the HL-LHC Era: A Community-Driven Technical Roadmap

IBERGRID 2026

Santiago de Compostela (Spain)

27-29 October 2026

PIC
port d'informació
científica

IFAE Ciemat RES RED ESPAÑOLA DE SUPERCOMPUTACIÓN España | digital 2026 Unió Europea Fons Europeu Next Generation GOBIERNO DE ESPAÑA Plan de Recuperación, Transformación y Resiliencia Next Generation Catalunya Generalitat de Catalunya Departament de Recerca i Universitats GOBIERNO DE ESPAÑA VICEPRESIDENCIA PRIMERA DEL GOBIERNO MINISTERIO DE ASUNTOS ECONÓMICOS Y TRANSFORMACIÓN DIGITAL SECRETARÍA DE ESTADO DE DIGITALIZACIÓN E INTELIGENCIA ARTIFICIAL

The Large Hadron Collider (LHC)

PIC
port d'informació científica

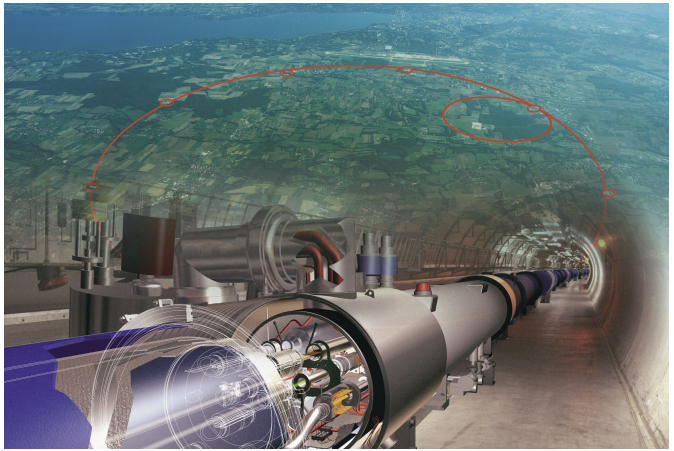
CERN • GENEVA, SWITZERLAND

World's Largest Particle Physics Laboratory

Exploring fundamental particles and forces in a 27-kilometer underground ring

Started Operations **~2010**

Engineering at Extreme Limits



CIRCUMFERENCE

27 km

Circular tunnel spanning across the border of France and Switzerland.

BEAM VELOCITY

99.99%

Speed of light reached by proton bunches inside vacuum pipes.

CRYOGENICS

1.9 K

Superfluid helium cooling superconducting magnets colder than outer space.

VACUUM

10^{-10} Torr

Twin beam pipes evacuated to prevent molecule collisions.

COLLISION RATE

40 MHz

40 million particle bunch crossings recorded every second.

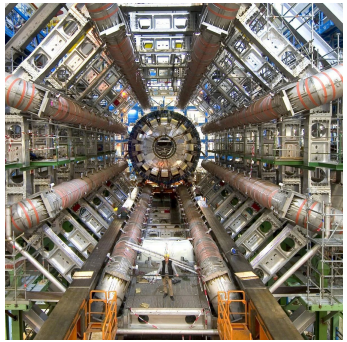
Four Main Particle Detectors

PIC
port d'informació
científica

ATLAS

A Toroidal LHC Apparatus

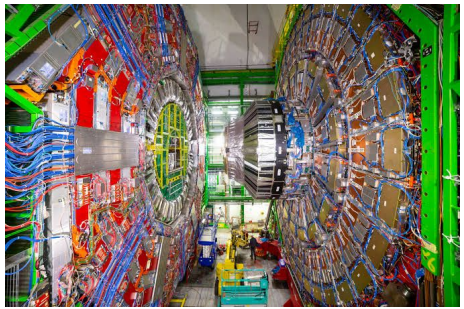
General-purpose detector for precision physics and searches beyond the Standard Model



CMS

Compact Muon Solenoid

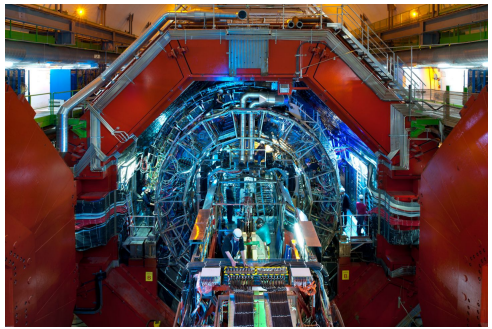
General-purpose detector for precision physics and searches beyond the Standard Model



ALICE

A Large Ion Collider Experiment

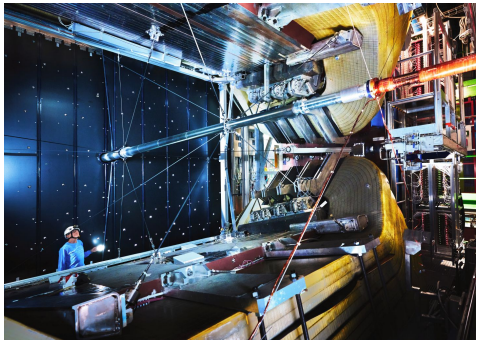
Heavy-ion detector studying quark-gluon plasma recreating early universe conditions.

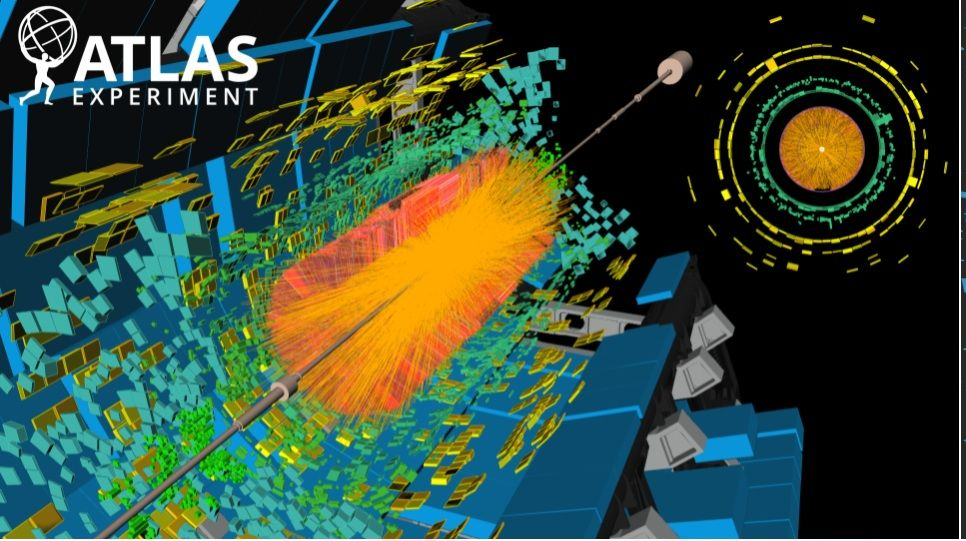


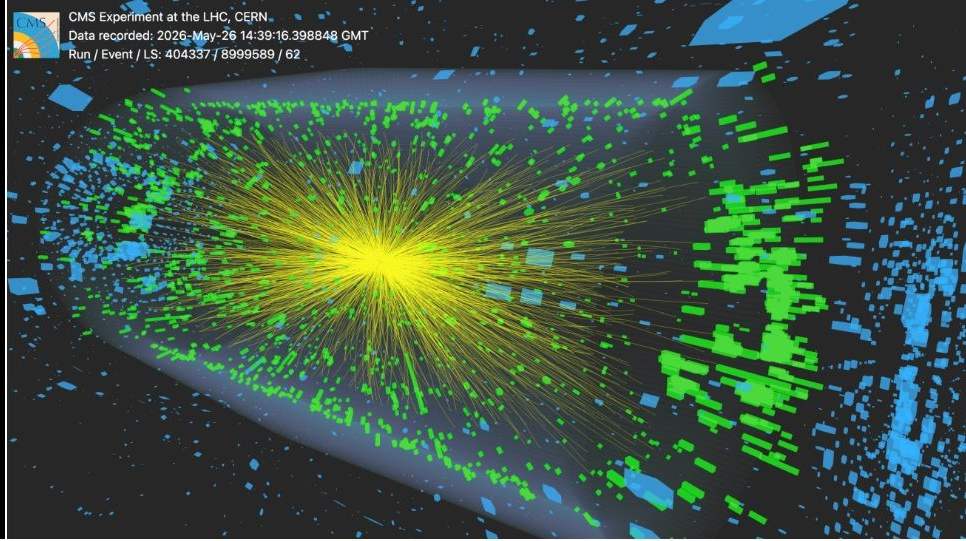
LHCb

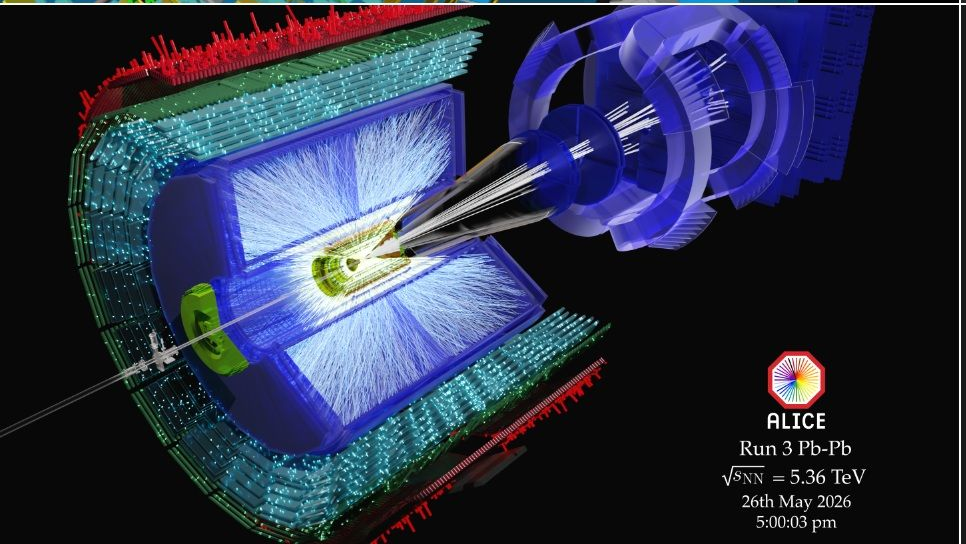
LHC-beauty Experiment

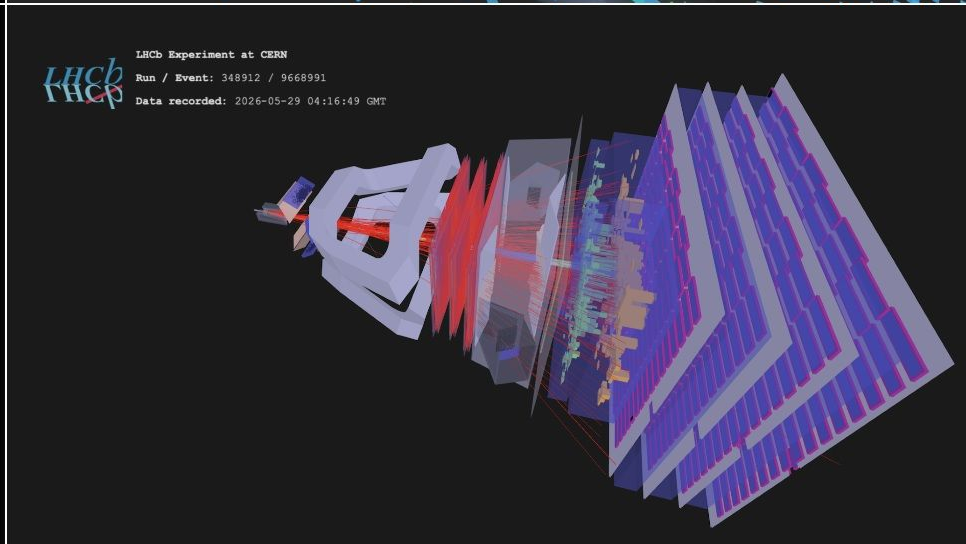
Precision detector studying beauty quarks to explain cosmic matter-antimatter imbalance.




ATLAS
EXPERIMENT


CMS Experiment at the LHC, CERN
Data recorded: 2026-May-26 14:39:16.398848 GMT
Run / Event / LS: 404337 / 8999589 / 62


ALICE
Run 3 Pb-Pb
$\sqrt{s_{NN}}$ = 5.36 TeV
26th May 2026
5:00:03 pm

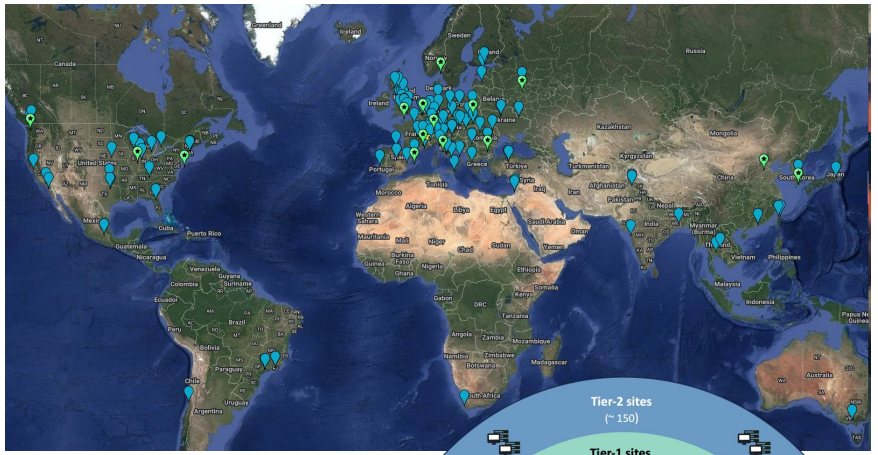

LHCb Experiment at CERN
Run / Event: 348912 / 9668991
Data recorded: 2026-05-29 04:16:49 GMT

The Worldwide LHC Computing Grid (WLCG)

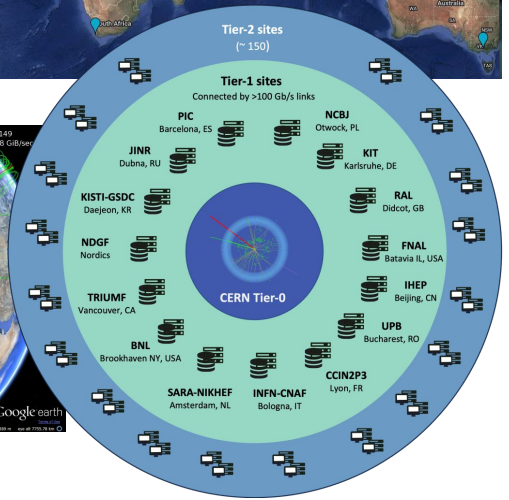
PIC
port d'informació científica

Distributed **high-throughput computing infrastructure** to store, process & analyze data produced by LHC experiments

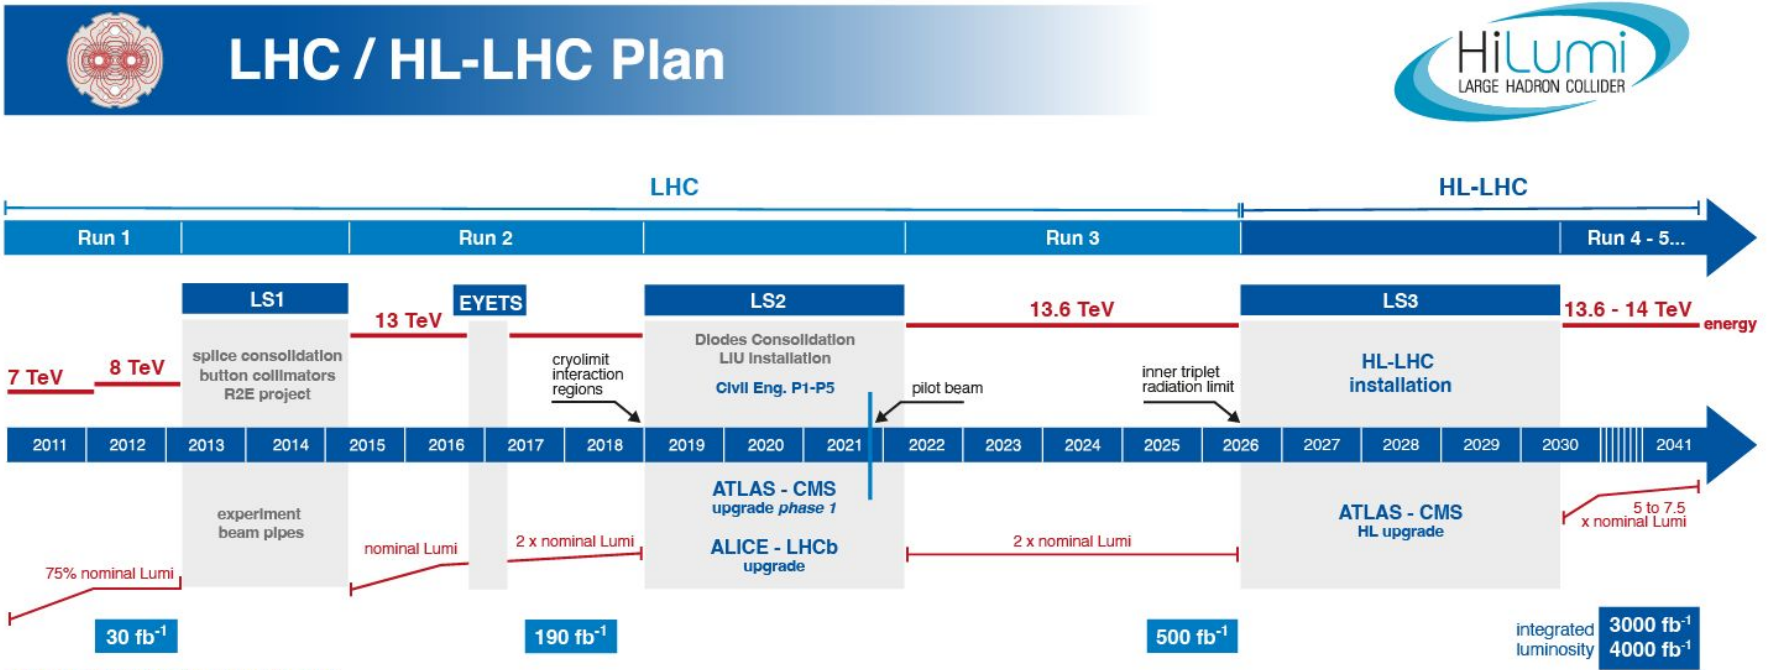
- 170 sites, 40+ countries
- >1 million CPU cores
- ~1.4 EB disk storage
- ~2.9 EB tape storage
- Optical private network (LHCOPN) and overlay over NRENs (LHCONE) with 100 Gbps links
- ~Tbps LAN bandwidth between compute and storage nodes at sites







WLCG: the future



Run 4 (HL/LHC, 2030+): **~10x more data**. **Revolutionary** changes required

WLCG Technical Evolution

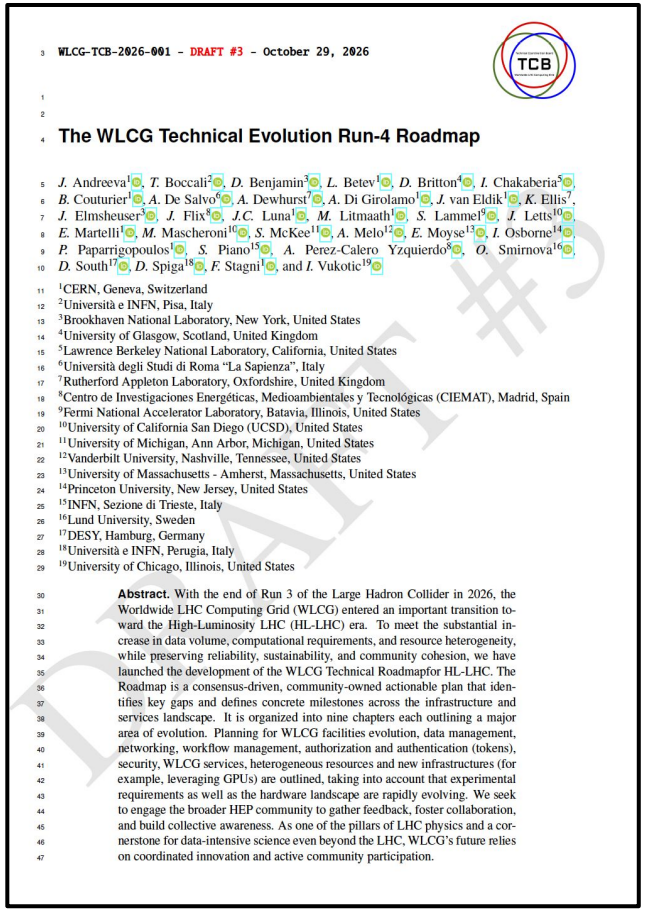
WLCG enables LHC scientific programs by providing essential computing infrastructure through a community-driven, "bottom-up" approach shaped including the computing facility collaborators themselves

WLCG **Technical Coordination** was established in Oct. 2024 to **coordinate this evolution**, **identify "gaps"** in scale or technology, and **outline a strategic technical roadmap** to bridge them

There are several fora where different aspects of the technical evolution are discussed, some in WLCG (e.g., the WLCG Open Technical Forum, operations coordination, DOMA project, various Task forces and working groups), others not (e.g., HEPiX Working Groups, IRIS-HEP, many others)

With this feedback, **we built a concrete, credible plan** to bring WLCG sites and services towards a shared goal (which will demonstrate to our reviewers e.g., LHCC, that we will be ready for Run 4)

WLCG Technical Evolution



WLCG-TCB-2026-001 - DRAFT #3 - October 29, 2026

The WLCG Technical Evolution Run-4 Roadmap

J. Andreeva[1], T. Boccali[2], D. Benjamin[3], L. Betev[1], D. Britton[4], I. Chakaberia[5], B. Couturier[1], A. De Salvo[6], A. Dewhurst[7], A. Di Girolamo[1], J. van Eldik[1], K. Ellis[7], J. Elmsheuser[3], J. Flix[8], J.C. Luna[1], M. Litmaath[1], S. Lammel[9], J. Letts[10], E. Martelli[1], M. Mascheroni[10], S. McKee[11], A. Melo[12], E. Moyse[13], I. Osborne[14], P. Paparrigopoulos[1], S. Piano[15], A. Perez-Calero Yzquierdo[8], O. Smirnova[16], D. South[17], D. Spiga[18], F. Stagni[1], and I. Vukotic[19]

[1]CERN, Geneva, Switzerland
[2]Università e INFN, Pisa, Italy
[3]Brookhaven National Laboratory, New York, United States
[4]University of Glasgow, Scotland, United Kingdom
[5]Lawrence Berkeley National Laboratory, California, United States
[6]Università degli Studi di Roma "La Sapienza", Italy
[7]Rutherford Appleton Laboratory, Oxfordshire, United Kingdom
[8]Centro de Investigaciones Energéticas, Medioambientales y Tecnológicas (CIEMAT), Madrid, Spain
[9]Fermi National Accelerator Laboratory, Batavia, Illinois, United States
[10]University of California San Diego (UCSD), United States
[11]University of Michigan, Ann Arbor, Michigan, United States
[12]Vanderbilt University, Nashville, Tennessee, United States
[13]University of Massachusetts - Amherst, Massachusetts, United States
[14]Princeton University, New Jersey, United States
[15]INFN, Sezione di Trieste, Italy
[16]Lund University, Sweden
[17]DESY, Hamburg, Germany
[18]Università e INFN, Perugia, Italy
[19]University of Chicago, Illinois, United States

Abstract. With the end of Run 3 of the Large Hadron Collider in 2026, the Worldwide LHC Computing Grid (WLCG) entered an important transition toward the High-Luminosity LHC (HL-LHC) era. To meet the substantial increase in data volume, computational requirements, and resource heterogeneity, while preserving reliability, sustainability, and community cohesion, we have launched the development of the WLCG Technical Roadmapfor HL-LHC. The Roadmap is a consensus-driven, community-owned actionable plan that identifies key gaps and defines concrete milestones across the infrastructure and services landscape. It is organized into nine chapters each outlining a major area of evolution. Planning for WLCG facilities evolution, data management, networking, workflow management, authorization and authentication (tokens), security, WLCG services, heterogeneous resources and new infrastructures (for example, leveraging GPUs) are outlined, taking into account that experimental requirements as well as the hardware landscape are rapidly evolving. We seek to engage the broader HEP community to gather feedback, foster collaboration, and build collective awareness. As one of the pillars of LHC physics and a cornerstone for data-intensive science even beyond the LHC, WLCG's future relies on coordinated innovation and active community participation.

DRAFT #3

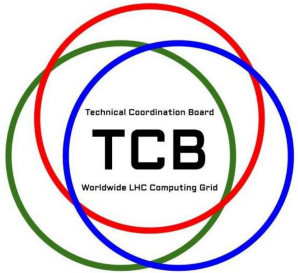
Draft #3 is **available for comments**

The document (65 pages) will be presented next week at the **WLCG/HSF 2026 Workshop** in Bologna, Italy

WLCG Technical Evolution

The WLCG Technical Roadmap is **organized thematically** into interconnected chapters:

Chapter 1 Experimental Requirements
Chapter 2 Facility Evolution
Chapter 3 Data Management
Chapter 4 Network Infrastructure and Management
Chapter 5 Workflow Management
Chapter 6 Authentication and Authorization Infrastructure (AAI)
Chapter 7 Cybersecurity
Chapter 8 WLCG Services
Chapter 9 New Architectures and Infrastructures
Chapter 10 WLCG and AI
Chapter 11 Sustainability



The Technical Roadmap is a bottom-up, consensus-driven plan bridging experiment needs with infrastructure capabilities

WLCG Technical Evolution

What is a Roadmap? It's a high-level **outline of a journey**, not just a destination

Like any journey it has:

- Milestones to guide the way
- Decision points - or forks in the road
- Can be fraught with risks that could have serious impacts
- And strategies to mitigate those risks



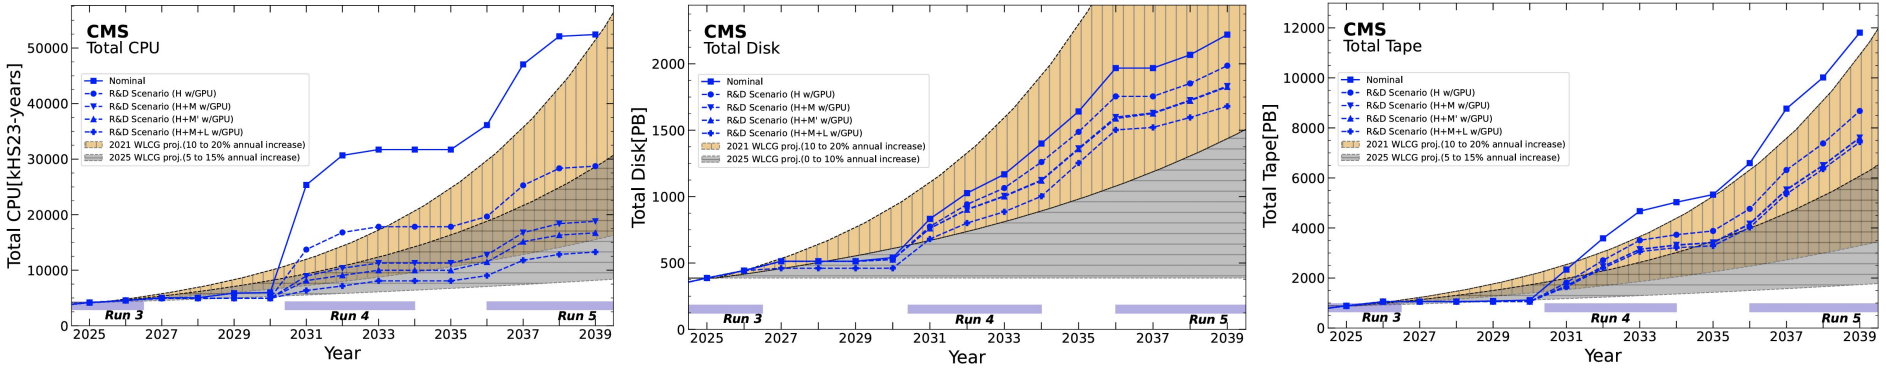
A first version of the WLCG Technical Roadmap has been made **public**

The document is expected to be **delivered to reviewers soon**

I will summarize the plan in each of the chapters in the rest of this presentation

Chapter 1 (Experimental Requirements)

Resource Projections: Depending on the success of experiment R&D, total Computing resource requirements in 2031* will be **2 to 2.5 times higher (optimistic)** or **3 to 4 times higher (pessimistic)** compared to 2026. Disk and Tape requirements are projected to **double**

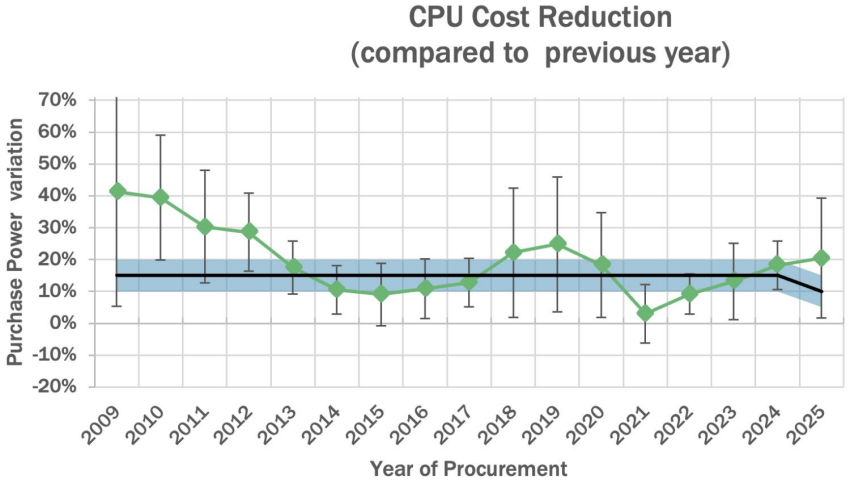


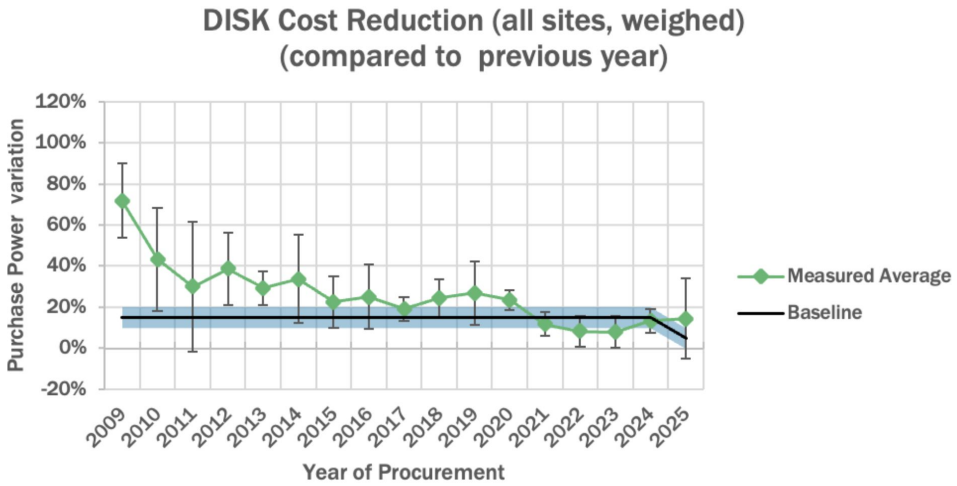
Figures from the CMS Offline Software and Computing for HL-LHC **Conceptual Design Report**

*2030 will mostly be a commissioning year with lower delivered luminosity. The full-scale production computing infrastructure needs to be deployed and ready by April 2031, which is considered the first full year of Run 4

Chapter 1 (Experimental Requirements)

Hardware Cost Crisis: The AI boom has shattered historical "flat-budget" assumptions, driving late-2025 price spikes of up to 4x for memory, 3x for SSDs, and 2x for HDDs





To navigate this fluid market, **WLCG will conduct yearly checkpoints with HEPiX** to monitor hardware market conditions alongside updated experiment requirements

Chapter 2 (Facility evolution)

Facilities face unprecedented challenges, including AI-driven **hardware price spikes**, storage capacity scaling faster than **read/write throughput**, and **environmental regulations** that might cap power delivery

Managing hardware expansion with flat operational budgets would require moving to **container-native orchestration** frameworks and utilizing **AI to ease operations**

Modern multi-core processors and accelerators demand **higher rack power densities**, necessitating transitions from air-cooling to liquid cooling or rear-door heat exchangers

Storage Consolidation: To reduce costs and human overhead, federations (like Germany and U.S. CMS) are concentrating mass storage at major centers while transforming Tier-2s into compute-heavy nodes backed by caches, relying on high-capacity WANs

Integrating non-owned **HPCs** and **Commercial Clouds** requires robust edge services. Additionally, **Analysis Facilities** handling interactive AI/ML workflows must be formally integrated for proper benchmarking, pledging, and accounting

Chapter 2 (Facility evolution)

Key Milestones include:

- **2026-2027:** Execute tape performance mini-challenges (Q4 2026), conduct a joint WLCG/HEPiX survey on site effort/power/AI (Q1 2027), and publish Tier-1 follow-up reports from DC27 (Q2 2027)
- **2028-2029:** Assess Tier-1 power, liquid cooling, and carbon compliance (Q4 2027), present storage capacity vs. throughput studies (Q2 2028), validate edge services (Q4 2028), and finalize operational validation from DC29 (Q2 2029)

Risk Mitigation: Facilities will address hardware market volatility through **HEPiX** technology tracking, possibly mitigating the storage I/O gap by deploying solid-state caching layers, and combat staffing deficits by adopting AI-driven operational tools

Chapter 3 (Data Management)

WLCG is scaling infrastructure via **DOMA Data Challenges**, moving from the 25% HL-LHC scale of DC24 to a 50% target for DC27 in Spring 2027. DC29 (2029) targets 100% scale

- DC27 will emphasize IPv6-only transfers, SciTags, jumbo frames, and AI-assisted operations

FTS is upgrading to **FTS4** to fix scheduling and database bottlenecks under heavy load, rolling out at CERN for DC27 and globally by Q3 2027

Upcoming **XRootD upgrades** will add native transfer monitoring, token authentication, and Tape REST API support

Rucio is preparing for exabyte-scale AI/ML workflows, improving HPC integration, and fully migrating to OIDC tokens (including self-minting to avoid issuer bottlenecks)

Key Risks: The primary threats are DC27 exposing unfixable infrastructure gaps and severe resourcing vulnerabilities due to small teams developing these critical tools

Chapter 4 (Network Infrastructure and mgt)

To meet HL-LHC demands, WLCG is targeting **WAN capacities of 400–800+ Gbps** for Tier-1 and large Tier-2 sites, utilizing a phased, experience-driven approach to balance costs and avoid bottlenecks

WLCG is aggressively **phasing out dual-stack environments** to achieve IPv6-only operations by 2030, alongside tuning systems to support 400–800 Gbps Network Interface Cards (NICs) and jumbo frames

Monitoring will be strengthened by deploying high-speed **perfSONAR** nodes and enabling **SciTags/Flow Marking** across 80% of sites by 2029 for precise data flow tracking

Operations will increasingly rely on **Software Defined Networking (SDN)** and tools like SENSE/Rucio for dynamic bandwidth allocation and automated path management

Proactive planning and capacity upgrades rely heavily on collaborations with global **Research and Education (R&E)** networks through forums like **LHCONE** and **LHCOPN**

Chapter 5 (Workflow Management)

Workflow systems are evolving beyond simple CPU slot allocation to manage **complex**, **heterogeneous resources** (such as GPUs and whole nodes) through advanced hardware matchmaking and co-scheduling

Systems must adapt to interface seamlessly with a **variety of backends**, including traditional Grid, HPC facilities, Clouds, and container-native infrastructures like Kubernetes

New Workflow Models: Workloads are shifting beyond traditional independent batch jobs to support multi-node tasks, interactive environments, and specialized AI/ML workflows (separating training from inference requirements)

WLCG will foster **collaboration to share interfaces**, resource descriptions, and integration WMS architecture strategies across different experiment specific systems

- E.g. CMS is adopting DiracX, developed together with LHCb

Initial steps focus on **community alignment**, starting with a WMs Community Workshop (Q1 2027) and the establishment of a dedicated WLCG WMs forum (Q2 2027)

Chapter 6 (AAI)

The Token Transition: WLCG is actively replacing legacy X.509 certificates and VOMS proxies with industry-standard JSON Web Tokens (OIDC)

To **minimize security risks** if leaked, tokens are designed with short lifetimes and tightly restricted scopes, particularly for operations that modify or delete storage

Core middleware (FTS, Rucio) and Compute Elements are **adapting to handle tokens** for job payloads and long-queued data transfers

To **prevent** central Identity and Access Management (IAM) services from becoming single **points of failure**, frameworks like Rucio are exploring "self-minted" storage tokens (similar to signed URLs) for high-rate data management

Key Milestones: Tokens for tape/ATLAS jobs (Q3 2026), Token Grand Challenge (Q4 2026), and a complete X.509 phase-out by 2028

Chapter 7 (CyberSecurity)

WLCG is embedding **cybersecurity** directly into all services, adapting to new capabilities like AI and analysis facilities

The newly established **WLCG Security Policies Working Group** will simplify and update policies, strengthen security expectations in the WLCG Memorandum of Understanding (MoU), and introduce self-assessment checklists for sites to measure their maturity

WLCG is creating a dedicated incident **playbook to handle** major compromises through clear escalation, identity blocking, and temporary site isolation

Fostering community learning through **anonymized incident reporting** that highlights security trends without compromising sensitive information

Launching a comprehensive **training roadmap** based on community needs to strengthen **hands-on skills in cybersecurity and Linux forensics**

Chapter 8 (WLCG services)

Accounting Modernization: Replacing the 20-year-old legacy APEL system with AUDITOR (for site-level aggregation) and PEAR (for the central repository) to accurately track heterogeneous resources, HPC, Cloud, and environmental metrics

CRIC (Information System) Evolution: Recently rebuilt in Python 3, CRIC is expanding to track detailed GPU and processor architectures, integrate storage space accounting, and support AI agents via a new MCP server

XRootD Transfer Monitoring: Closing the critical visibility gap exposed during DC24 by rolling out dedicated site-level collectors and integrating dCache billing records to achieve complete throughput monitoring

Key Milestones: Upcoming targets include tracking GPUs in CRIC (Q2 2027), launching the central PEAR accounting repository (Q4 2027), and deploying XRootD monitoring across all WLCG sites by the start of Run 4

Chapter 9 (New Architectures and Infra.)

Embracing Heterogeneity: WLCG is moving beyond x86 architectures by integrating ARM (already in production for ALICE and ATLAS), preparing for GPU-heavy workflows, and running "technology watches" on emerging hardware like RISC-V, FPGAs, and NPUs

Benchmarking & Accounting: Developing HEPScore4GPU with HEPIX (which will evolve into HEPScore2X) to evaluate heterogeneous nodes, alongside new accounting models to accurately track total-cost-of-ownership and GPU-hours

Integrating Non-Owned Resources: Adapting provisioning and federated identity tools to seamlessly utilize external capacity from HPC centers, Commercial Clouds, Analysis Facilities, and emerging EuroHPC AI Factories

Key Milestones: Upcoming targets include deploying GPU benchmarking (Q3 2028), establishing GPU accounting (Q4 2028), and executing a full heterogeneous infrastructure integration test (Q4 2028)

Chapter 10 (WLCG and AI)

WLCG services must be redesigned to be operable by AI agents, establishing CRIC as the single source of truth for configuration and adding programmatic access via APIs and Model Context Protocol (MCP) interfaces

Operational and monitoring **data must stream in real-time** to indexed stores, allowing AI agents to immediately correlate events across services instead of relying on delayed, human-centric dashboards

DC27 will deploy **AI assistants** (Archi and AEGIS) for system diagnosis and human-approved actions, with the goal of transitioning to fully pre-authorized, autonomous agent actions by DC29

To prevent fragmented development, WLCG will consolidate the diverse AI requirements of all experiments into a **single, quarterly-updated document to guide** central service upgrades

Securing AI Agents: Mitigating write-access risks through strictly scoped tokens, audit logging, and a phased rollout of agent autonomy

Chapter 11 (Sustainability)

Environmental Metrics: Prioritizing reductions in power usage and carbon footprints by introducing and tracking new "energy-to-solution" metrics (e.g., Joules per event) alongside traditional time-to-solution

Driving Efficiency: Recognizing that the largest sustainability gains come from software optimizations (yielding >10x improvements), followed by the adoption of energy-efficient hardware and architectures

Operational Sustainability: Leveraging Agentic AI and operational intelligence to automate routine tasks, ensuring WLCG can maintain operations effectively despite flat or declining personnel effort

Summary

The **WLCG Technical Roadmap** delivers a consensus-driven, actionable plan to bridge our technical gaps for the HL-LHC era

- The **document** will be presented at next week **WLCG Workshop 2026**
- The document will be soon reviewed by the LHCC
- The milestones completion and further discussions will be driven by the WLCG TCB

While we have **detailed timelines** for facilities, data, and services, many complex challenges still r**equire significant work and innovation**

WLCG's future relies entirely on **coordinated innovation** and **active community participation**

Call to Action: we need the community to actively participate in this process and help us engineer the (successful) future of WLCG

Acknowledgements

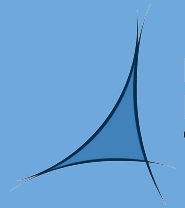
We thank all of our colleagues for their contributions to the WLCG Technical Roadmap, Open Technical Forum Meetings, and the Technical Collaboration Board, including our partners in HEPiX, throughout the WLCG, in the experiments, software and middleware development teams, and at the facilities … a list too long to write here, but colleagues without whose work none of this would have been possible

The WLCG Technical Evolution Run-4 Roadmap

Authors (The WLCG Technical Coordination Board and ex-officio members)

J. Andreeva[1], *T.* Boccali[2], *D.* Benjamin[3], *L.* Betev[1], *D.* Britton[4], *I.* Chakaberia[5], *B.* Couturier[1], *A.* De Salvo[6], *A.* Dewhurst[7], *A.* Di Girolamo[1], *J.* van Eldik[1], *K.* Ellis[7], *J.* Elmsheuser[3], *J.* Flix[8], *J.C.* Luna[1], *M.* Litmaath[1], *S.* Lammel[9], *J.* Letts[10], *E.* Martelli[1], *M.* Mascheroni[10], *S.* McKee[11], *A.* Melo[12], *E.* Moyse[13], *I.* Osborne[14], *P.* Paparrigopoulos[1], *S.* Piano[15], *A.* Perez-Calero Yzquierdo[8], *O.* Smirnova[16], *D.* South[17], *D.* Spiga[18], *F.* Stagni[1], and *I.* Vukotic[19]

Thank you!



PIC
port d'informació
científica